

Whose Opinion to follow in Multihypothesis Social Learning? A Large Deviations Perspective

Wee Peng Tay, *Member, IEEE*

Abstract—We consider a multihypothesis social learning problem in which an agent has access to a set of private observations and chooses an opinion from a set of experts to incorporate into its final decision. To model individual biases, we allow the agent and experts to have general loss functions and possibly different decision spaces. We characterize the loss exponents of both the agent and experts, and provide an asymptotically optimal method for the agent to choose the best expert to follow. We show that up to asymptotic equivalence, the worst loss exponent for the agent is achieved when it adopts the 0-1 loss function, which assigns a loss of 0 if the true hypothesis is declared and a loss of 1 otherwise. We introduce the concept of hypothesis-loss neutrality, and show that if the agent adopts a particular policy that is hypothesis-loss neutral, then it ignores all experts whose decision spaces are smaller than its own. On the other hand, if experts have the same decision space as the agent, then choosing an expert with the same loss function as itself is not necessarily optimal for the agent, which is somewhat counter-intuitive. We derive sufficient conditions for when it is optimal for the agent with 0-1 loss function to choose an expert with the same loss function.

Index Terms—Social learning, decentralized detection, error exponent, social network, Internet of Things.

I. INTRODUCTION

In an increasingly connected world, our opinions on a phenomenon of interest or event are often not only influenced by our direct independent observations, but also by other people's public opinions on related events. In an online social network like Twitter or Facebook, users' opinions and postings are often influenced by the opinions of those they are connected to or are "following" in the social network [1]–[4]. For example, in viral marketing using social networks, marketing companies often target a few influential nodes in the network to help them push a product [5], [6]. Similarly, widespread online access has made it easier for us to follow the opinions of experts like celebrities and industry insiders, who may have access to private information that we are unaware of. For example, we may be interested to determine the financial health of a publicly listed company. In addition to our own observations about the company through its annual financial reports and stock prices, we may also choose to incorporate the "expert" opinion provided by financial blogs like [7] and [8]. In all these examples, inference about a phenomenon of interest is

not only based on direct observations but also the opinions of other entities. This is known as *social learning* [9], [10].

A further example is the Internet of Things (IoT) framework [11], [12]. Sensors each make their own private observations but collaborate by exchanging public information. This can be viewed as a "physical social network" of devices, cooperating to improve their situational awareness. Sensors originally designed for a specific purpose may collaborate with other sensors to perform inference on a phenomenon they were not specifically designed for. However, in order to ensure energy efficiency, each sensor needs to intelligently choose which other sensors to collaborate with since not all sensors may provide information relevant to it. For example, a sensor trying to estimate the temperature in a particular room of a building may choose to incorporate information from other temperature sensors or sensors tracking the number of occupants in the building, instead of information from a vibration monitoring sensor.

In this paper, we investigate the problem of multihypothesis social learning in which an agent can select an expert opinion from a group of experts, to incorporate into its final decision in order to minimize its expected loss. Specifically, we consider an agent 0 who wishes to choose from a set of M hypotheses based on its own observations as well as the opinion of an expert chosen from a set of K possible experts. Each expert has access to a set of private observations, which they use to form their own opinions about the true hypothesis, subject to their own local loss functions or biases, which may differ from that of agent 0. By taking into account the experts' individual biases, our goal is to find an asymptotically optimal expert choice in order to minimize agent 0's expected loss.

A. Related Work

The problem of selecting the best expert opinion to follow is related to the problem of *decentralized detection* or decentralized hypothesis testing, which has been extensively studied in [13]–[23] and the references therein. In the decentralized detection problem, each agent has its own private observations, but cooperate with each other so that the whole network of agents reaches a decision regarding a common underlying phenomenon of interest. Here, which agent passes information to which other agent is determined by a known network topology like the parallel configuration [24], tandem network [25]–[27], and tree architectures [19]–[22]. Therefore, all these works do not consider the problem of selecting which other agent's opinion to follow. Furthermore, in a decentralized detection problem, all agents are assumed to have the same hypotheses

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

This research is supported by the MOE Tier 2 grants MOE2013-T2-2-006 and MOE2014-T2-1-028. The author is with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. E-mail: wptay@ntu.edu.sg

and loss functions for declaring the wrong hypothesis, or have the common goal of minimizing the loss at a last agent known as the fusion center.

We study a related but somewhat different problem from the decentralized detection problem. We consider the scenario where experts may make decisions according to their own biases, instead of minimizing the loss of a particular agent or fusion center. In [9], the authors consider binary hypothesis testing in a social network, in which agents sequentially observe the opinions of a stochastically generated neighborhood of agents in the network. Each agent has the same 0-1 loss function,¹ and conditions are derived for the Bayesian error probability of the n th agent to approach to zero as n becomes large. The network model we consider in this paper is equivalent to a two-layer hierarchical tree network, which is much simpler compared to that studied in [9]. This is because our goal is to analyze how an expert's opinion impacts that of a particular agent. In addition, we consider a multihypothesis testing problem in which each agent has different loss functions or even different number of hypotheses.

Since finding optimal decision rules for general decentralized hypothesis testing problems is NP-complete [28], it is difficult to find analytical characterizations for the best expert choice in networks of moderate size, and where agents have private observations that are not conditionally independent given the underlying hypothesis. Therefore, most of the literature in decentralized detection has focused on the case where observations are conditionally independent [13]–[22]. The same challenge applies to the social learning problem in this paper, and we therefore consider only the case where agents' private observations are conditionally independent. Although the problem becomes numerically tractable under this assumption, it remains hard to characterize the agents' optimal policies and decision rules analytically (see [21] for a discussion). To overcome this difficulty, we consider the regime where each agent has access to an asymptotically large number of private observations (e.g., over a sufficiently long period of time) but each agent's observation sources and loss functions may differ. This allows us to adopt a large deviations perspective to the expert choice problem, and derive analytical characterizations of the optimal choice and agent strategies.

The reference [29] considers a multihypothesis testing problem in a parallel configuration with a single agent, which is allowed to take actions that affect the observations it makes. The distribution of the observation at each time step depends on the action of the agent in the previous time step, and has bounded Kullback-Liebler divergences under any pair of hypotheses. The agent's aim is to minimize the error exponent corresponding to the 0-1 loss function. Our work can be viewed as a generalization of a result in [29] (which also considers the sequential detection problem that we do not study here) to the case where one of the agent observations is the opinion of an expert, which itself has asymptotically many private observations leading to unbounded Kullback-

Liebler divergences for its opinion. Furthermore, we consider general loss functions and the corresponding loss exponent, which is a generalization of the error exponent in [29]. The characterization of the error exponents corresponding to 0-1 loss functions in multihypothesis testing has also been studied in [30], [31], which derive the achievable error exponents region.

B. Our Contributions

Our goal is to optimize the loss exponent of agent 0, the absolute value of which is the rate of decay of the expected loss incurred by agent 0 as the number of private observations grows large. We suppose that agent 0 has a decision space consisting of M hypotheses, and has a choice of K expert opinions to follow.

- (i) We consider a general M -ary multihypothesis social learning framework in which an agent 0 and every expert have loss functions that depend on the number of private observations the agent or expert has access to. Each expert makes a decision that takes values from a decision space, the size of which may not be the same as M . We characterize the loss exponent of each expert, and provide an asymptotically optimal policy for an expert to achieve its optimal loss exponent (Theorem 2).
- (ii) We derive the optimal loss exponent for agent 0 after it has incorporated the opinion of a particular expert, and provide an asymptotically optimal method for agent 0 to choose the best expert to follow (Theorem 5). We also show that the worst loss exponent, up to asymptotic equivalence,² for agent 0 is achieved when agent 0 adopts the 0-1 loss function (cf. Remark 3 of Section IV-A).
- (iii) We introduce the concept of hypothesis-loss neutrality in Section IV-B, and show that if agent 0 adopts a hypothesis-loss neutral policy, then the opinion of any expert who has a decision space with number of states strictly less than M , is ignored by agent 0 (Proposition 6). In this case, additional information is useless, which is somewhat unexpected.
- (iv) We show that if all experts have the same decision space as agent 0, then it is not necessarily optimal for agent 0 to choose the expert with the same loss function as itself. This is surprising as conventional wisdom seems to suggest otherwise. If agent 0 adopts the 0-1 loss function, we derive sufficient conditions for when it is optimal to choose an expert who also utilizes the 0-1 loss function (Proposition 7).

In this work, we do not address the case where agent 0's opinion may be utilized by one of the experts, which result in much more complex opinion dynamics than that considered in this paper. This leads to the interesting question of whether there exists an equilibrium in the choice of loss functions in a network of agents who can incorporate opinions from each other. This is however out of the scope of the current work, and will be addressed in our future research.

The rest of this paper is organized as follows. In Section II, we describe our system model, problem formulation and

¹The 0-1 loss function assigns a loss of 0 if the agent declares the same hypothesis as the true underlying hypothesis, and a loss of 1 otherwise.

²See Section II-C for the definition of asymptotic equivalence.

assumptions. We characterize the agents' loss exponents in Section III, and provide asymptotically optimal policies to achieve the optimal loss exponents. In Section IV, we discuss the optimal expert choice for agent 0, and derive insights into this choice by making simplifying assumptions. We conclude in Section V. Appendix A contains a brief review of some basic definitions of large deviations theory, and we defer all proofs to Appendix B. Appendix C contains a characterization of the asymptotic decision regions, which are introduced in Section III-A.

II. PROBLEM FORMULATION

In this section, we define our system model, assumptions and some notations. We consider an underlying measurable space (Ω, \mathcal{F}) on which all random variables in this paper are defined. We adopt the following notations throughout this paper. Let \mathbb{R} be the space of real numbers, and let $\langle t, z \rangle$ be the inner product of the vectors t and z in \mathbb{R}^{M-1} . For $z \in \mathbb{R}^{M-1}$, let $z^0 = (0, z) \in \mathbb{R}^M$ be the vector augmented with a zero as the first element. The range of integers $a, a+1, \dots, b$ is denoted as $[a, b]$. The notations $x[1 : n]$ and $(x_i)_{i=1}^n$ are used to represent the sequences $x = (x[1], x[2], \dots, x[n])$ and (x_1, x_2, \dots, x_n) , respectively. We also make use of $x[i]$ to denote the i th element in the vector or sequence x .

A. Learning From an Expert

Suppose that an agent 0 wishes to determine the underlying hypothesis H associated with a phenomenon of interest, which from agent 0's frame of reference can be modeled by a set of probability measures $\{\mathbb{P}_m : m = 0, 1, \dots, M-1\}$ on the space (Ω, \mathcal{F}) . Let \mathbb{E}_m be the mathematical expectation under \mathbb{P}_m , and let $H = m \in [0, M-1]$ if all agents' private observations have distributions derived from the probability measure \mathbb{P}_m . We suppose that $H = m$ has prior probability $\pi_m \in (0, 1)$. The agent 0 can choose to incorporate the opinion of an expert, chosen from a set of expert agents $\{1, 2, \dots, K\}$ (see Figure 1). If there is no need to distinguish between agent 0 and the experts, we use the generic term "agent" to refer to either of them.

Each agent $k \in \{0, \dots, K\}$ has access to a set of private observations $Y_k[1 : n_k] = (Y_k[1], \dots, Y_k[n_k])$. Conditioned on $H = m$, we assume that each $Y_k[l]$, for $l = 1, \dots, n_k$, is drawn independently from a conditional distribution belonging to the set $\{\mathbb{P}_m^\gamma : \gamma \in \Gamma_k\}$, where Γ_k is a finite index set representing the information sources of agent k . The distribution set $\{\mathbb{P}_m^\gamma : \gamma \in \Gamma_k\}$ varies from agent to agent, and models the differences in quality of information that each agent may have access to. We assume that for each $\gamma \in \Gamma_k$, the probability measures \mathbb{P}_m^γ for $m = 0, \dots, M-1$, are absolutely continuous with respect to (w.r.t.) each other, and are known to all the agents. We further assume that conditioned on H , all the random variables $(Y_k[1 : n_k])_{k=0}^K$ are independent.³

To motivate our setup, consider the example alluded to in Section I, where an agent 0 is interested to determine the




FIG. 1: Learning network in which agent 0 incorporates the decision of an agent $k \in \{1, 2, \dots, K\}$.

financial health of a large publicly listed conglomerate. In addition to its own observations $Y_0[1 : n_0]$ about the company through annual financial reports and other publicly available indicators like stock prices, the agent 0 may also choose to incorporate the expert opinion of an investment analyst. Because of limited financial resources (since most analyst reports are not free), agent 0 can only choose to subscribe to one analyst, and has to make an optimal choice of which analyst to use. In most of this paper, for simplicity, we restrict ourselves to the case where the opinion of a single expert is considered in the decision making of agent 0, and show how to generalize our results to a finite set of experts later in Remark 4 of Section IV. Note also that the number of information sources for each agent is also typically finite as in this example, which justifies our assumption that $|\Gamma_k| < \infty$ for all $k \geq 0$.

Let $x_k = (x_k[\gamma])_{\gamma \in \Gamma_k}$ be a vector of non-negative weights summing to one, where $x_k[\gamma]$ is the number of private observations of agent k that have conditional distribution \mathbb{P}_m^γ when the hypothesis $H = m$. We let the remaining $n_k - \sum_{\gamma \in \Gamma_k} x_k[\gamma] n_k$ observations be drawn from an arbitrary distribution from the set $\{\mathbb{P}_m^\gamma : \gamma \in \Gamma_k\}$. Since we are concerned about asymptotics in this paper, the choice of the arbitrary distribution is immaterial to our analysis. We say that x_k is a *policy* for agent k . An agent chooses its policy in order to minimize its expected loss, as described below. Let $\mathbb{S}(\Gamma_k)$ be the simplex consisting of all agent policies.

Based on its observations, each expert $k \geq 1$ makes a decision $D_k = \gamma_k(Y_k[1 : n_k]) \in [0, d_k - 1]$ by minimizing a local loss criterion. We assume that every expert $k \geq 1$ chooses its policy x_k and decision rule γ_k so that

$$\mathbb{E}[C_k(H, D_k, n_k)] \quad (1)$$

is minimized, where $C_k(m, d, n_k)$ is a non-negative finite loss incurred if the decision of expert k is d when the true hypothesis is $H = m$. We can think of $C_k(\cdot, \cdot, \cdot)$ as encoding the "bias" of expert k . We call this the *loss function* of expert k , and

$$\limsup_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, n_k)] \quad (2)$$

the *loss exponent* of expert k . An example of a loss function is the 0-1 loss function: $C_k(m, m, n_k) = 0$, and $C_k(m, d, n_k) = 1$ for all $n_k \geq 1$ if $d \neq m$, which results in (1) being the error probability considered in [32]. In this paper, we consider

³Although two experts may have the same information source, we assume their private observations from this source are independent due to their own noisy interpretations of the same piece of information.

the regime where experts have asymptotically large number n_k of private observations, which corresponds to the case where experts are very experienced in their respective fields (as determined by Γ_k).

By allowing the loss function to depend on the number of private observations, we can model various practical applications. For example, the expert k may itself be part of a decentralized detection network like a tree configuration [21], [26], in which case its goal is to minimize the Bayesian error probability at a fusion center. Then, its loss function decays exponentially fast in the number of private observations n_k . In the publicly listed conglomerate example, suppose that one of the hypotheses corresponds to the conglomerate being financially bankrupt. The loss associated with a missed detection of this hypothesis can be modeled to be exponentially larger than the loss associated with a missed detection of another more benign hypothesis.

In our model, expert k 's decision space consists of d_k states, where d_k may not equal to M , the number of hypotheses that agent 0 is interested in. This allows us to model scenarios where different agents may have different models for the underlying state of the world. Using the conglomerate example described above, agent 0 may be interested in the financial health of the whole conglomerate, while a particular analyst may only be interested in the real estate arm of the conglomerate. Nevertheless, the performance of the real estate arm has a bearing on the overall health of the conglomerate. In particular, expert k may not distinguish between two hypotheses, say $H = M-2$ and $H = M-1$. This can be modeled by assuming that the expert declares the decision $d = m$ when $H = m < M-2$ and $d = M-2$ if $H = M-2$ or $M-1$. In this case, we let $C_k(M-2, d, n_k) = C_k(M-1, d, n_k) = 1$ for all $d < M-2$ and $C_k(M-2, M-2, n_k) = C_k(M-1, M-2, n_k) = 0$. This model is also general enough to model different decision spaces that may not map directly to any of the hypotheses of agent 0.

On the other hand, we note that without loss of generality, there is no need to consider the case where an expert uses a model in which there are more than M hypotheses, since agent 0 can always choose a M that is sufficiently large in its model. In addition, since our analysis is based on the frame of reference of agent 0, it has no knowledge of any additional hypotheses, and in practical applications, it simply assumes that expert k minimizes the expected loss given in (1). We assume that based on the publicized expertise of each expert, agent 0 knows its loss function decay rates, defined in Assumption 1 below.

B. Loss of Agent 0

Let $D_0(k) = \gamma_0(Y_0[1 : n_0], D_k) \in [0, M-1]$ be the decision made by agent 0 after incorporating the opinion D_k of agent k . The expected loss of agent 0 is

$$\mathbb{E}[C_0(H, D_0(k), n_0)],$$

where $C_0(m, d, n_0)$ is the non-negative loss incurred by agent 0 if it decides in favor of hypothesis d when the true hypothesis is $H = m$, and the number of private observations is n_0 . We

make the following assumptions regarding the loss function of each agent.

Assumption 1.

- (i) For all $n_0 \geq 1$ and $m \in [0, M-1]$, we have $C_0(m, m, n_0) = 0$, and for every $d \in [0, M-1]$ such that $d \neq m$, we have $0 < C_0(m, d, n_0) \leq 1$ with

$$-\lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log C_0(m, d, n_0) = c_0(m, d) \equiv c_0(m) \in [0, \infty).$$

- (ii) For each agent $k \geq 1$, we have for every $m \in [0, M-1]$ and $d \in [0, d_k-1]$,

$$-\lim_{n_k \rightarrow \infty} \frac{1}{n_k} \log C_k(m, d, n_k) = c_k(m, d) \geq 0.$$

We assume that agent 0 knows the loss function decay rates $c_k(\cdot, \cdot)$, for all $k = 1, \dots, K$, based on the publicized expertise of each expert.

Since the number of agents and decision states are finite, we can normalize $C_k(m, d, n_k)$ by $\max_{m,d,k} C_k(m, d, n_k)$ or $\sum_{d,k} C_k(m, d, n_k)$ so that there is no loss in generality in assuming that $C_k(m, d, n_k) \leq 1$ for all m, d , and k . In Assumption 1(i), we make the simplifying assumption that the loss function decay rate for agent 0 w.r.t. a particular hypothesis is the same for all wrong decisions. This is because otherwise, easy examples can be constructed in which agent 0 declares the wrong hypothesis with high probability for large n_0 . To see this, suppose that all hypotheses have the same prior probability, $H = 0$ is the true hypothesis, and agent 0's observations are all drawn from the same conditional distribution. The loss of agent 0 if it makes the decision d is, with high probability, approximately proportional to

$$\sum_{m \neq d} e^{-n_0(c_0(m,d) + K_{0m})},$$

where K_{0m} is the Kullback-Liebler divergence of the conditional distribution under $H = 0$ versus that under $H = m$ [33]. If $c_0(0, 1) > c_0(1, 0) + K_{01}$, while $c_0(m, 1) > c_0(m, 0)$ for all $m \neq 0, 1$, then agent 0 decides in favor of $H = 1$ instead of $H = 0$ when n_0 is large. This is clearly an undesirable model as $c_0(m, 1) > c_0(m, 0)$ for all $m \neq 0, 1$ implies that the agent imposes an exponentially smaller loss when declaring $H = 1$ versus $H = 0$. In this case, an arguably more appropriate modeling approach is to merge the hypothesis $H = 1$ into another hypothesis.

In Assumption 1(ii), we assume that agent 0 knows the loss function decay rates of the experts. This assumption may not hold in some practical applications, in which case we need to impose uncertainties on agent 0's knowledge about the experts, similar in spirit to [34]. This unfortunately makes the problem much more challenging, and is out of the scope of our current work. The results in this work serve as the basic foundation on which the more challenging problem in which agent 0 has limited knowledge of the experts, can be addressed in future research.

C. Asymptotic Equivalence and Optimal Loss Exponent

Consider two loss functions with loss decay rates $c_k(m, d)$ and $c'_k(m, d)$ respectively. If $c_k(m, d)$ and $c'_k(m, d)$ differ by the same constant for all m and d , then their corresponding loss exponents (2) are different but the optimal policies to minimize (2) are the same. We therefore say that two loss functions are *asymptotically equivalent* if their respective loss decay rates $c_k(m, d)$ and $c'_k(m, d)$ differ by the same constant for all m and d . It can be shown that each loss function belongs to an equivalence class, with the equivalence relation being asymptotic equivalence. We call a loss function with $\min_{m,d} c_k(m, d) = 0$ a *canonical loss function* of its equivalence class. For example, the 0-1 loss function is a canonical loss function of its equivalence class. In another example, consider the set of all loss functions $C_k(m, d, n_k)$ with $\sum_d C_k(m, d, n_k) = C$ for each m and n_k , which imposes a total loss of C for each hypothesis. Then, it can be shown that each of these loss functions is a canonical loss function.

We say that two loss exponents are equivalent if they have the same policy and their loss functions are asymptotically equivalent. For fair comparison of loss exponents, we will always assume canonical loss functions. We are interested to characterize the optimal loss exponent⁴

$$\min_{1 \leq k \leq K} \limsup_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{E} [C_0(H, D_0(k), n_0)], \quad (3)$$

when the number of private observations of agent 0 becomes large. Note that although agent 0 is allowed an asymptotically large number of private observations, the quality of information available to agent 0 is constrained by the set Γ_0 . The experts' information sources $\{\Gamma_k : k = 1, \dots, K\}$ are different from Γ_0 , and may thus improve the loss exponent of agent 0. We assume that for all experts $k \geq 1$, $q_k = \lim_{n_0 \rightarrow \infty} n_k/n_0$ exists as a limit. If the number n_k of private observations of expert k is such that $q_k = 0$, agent 0 will ignore the opinion of expert k as its opinion becomes asymptotically negligible compared to agent 0's private observations. Therefore, without loss of generality, we assume that $q_k > 0$ for all $k \geq 1$.

D. Technical Definitions

In this subsection, we define notations that will be commonly used throughout the paper. We rely heavily on the mechanisms of large deviations theory to characterize the agents' loss exponents (2) and (3). We refer the reader to Appendix A for a brief overview of some basic concepts, and to [35] for a detailed treatment of large deviations theory. As the theory of large deviations in hypothesis testing problems utilizes log moment generating functions of log likelihood ratios, and their Fenchel-Legendre transforms in order to characterize the loss exponents, we define these necessary quantities in the following.

For any given random variable X with marginal distribution \mathbb{P}_i^X under hypothesis $H = i$, we abuse notation by letting $\ell_{ij}(X)$ be the Radon-Nikodym derivative (or likelihood ratio) of \mathbb{P}_i^X w.r.t. \mathbb{P}_j^X . Note that $\ell_{ij}(X)$ is a random variable that

⁴Note that because of Assumption 1, the loss exponent is *negative*, with a more negative loss exponent corresponding to a faster loss decay rate.

depends on the distributions and realization of X . By convention, we let $\ell_{ii}(X) = 1$ for all values of i and all realizations of X . In addition, for simplicity, we let $\ell_{ij}^\gamma = d\mathbb{P}_i^\gamma/d\mathbb{P}_j^\gamma$ be the Radon-Nikodym derivative of \mathbb{P}_i^γ w.r.t. \mathbb{P}_j^γ , for each $\gamma \in \bigcup_k \Gamma_k$.

Let $Z^\gamma = (\log \ell_{m0}^\gamma)_{m=1}^{M-1}$ be a vector of log likelihood ratios, and the log moment generating function of Z^γ under $H = m$ be

$$\xi_m(\gamma, t) = \log \mathbb{E}_m [\exp(\langle t, Z^\gamma \rangle)],$$

for all $t \in \mathbb{R}^{M-1}$. For a policy $x = (x[\gamma])_{\gamma \in \Gamma_k}$ of an agent k , the weighted log moment generating function is then given by

$$\varphi_m(t, x) = \sum_{\gamma \in \Gamma_k} x[\gamma] \xi_m(\gamma, t), \quad (4)$$

and its Fenchel-Legendre transform [35] is

$$\Phi_m^*(z, x) = \sup_{t \in \mathbb{R}^{M-1}} \{ \langle t, z \rangle - \varphi_m(t, x) \}, \quad (5)$$

where $z \in \mathbb{R}^{M-1}$. The function $\Phi_m^*(z, x)$ will be shown in Theorem 2 to characterize the rates of decay of the probabilities $\mathbb{P}_m(D_k = d)$, similar to the rate functions in large deviations theory [35].

When an expert k 's decision space has size $d_k = M$, the rates of decay of the probabilities $\mathbb{P}_m(D_k = d)$ become easier to characterize since we can now interpret a decision d of the expert k to be in favor of hypothesis d . Then, under the true hypothesis $H = m$, the log moment generating function of interest is the one involving the likelihood ratios of the observation distributions under hypothesis d versus hypothesis m , given by

$$\Lambda_{ij}(s, x) = \sum_{\gamma \in \Gamma_k} x[\gamma] \log \mathbb{E}_i [(\ell_{ji}^\gamma)^s]. \quad (6)$$

The Fenchel-Legendre transform of Λ_{ij} is then given by

$$\Lambda_{ij}^*(z, x) = \sup_{s \in \mathbb{R}} \{ sz - \Lambda_{ij}(s, x) \},$$

where $z \in \mathbb{R}$.

We assume that the distributions $\{\mathbb{P}_m^\gamma, \gamma \in \bigcup_k \Gamma_k\}$ are well-behaved for all $m \in [0, M-1]$ in the following assumption, which holds for example in the case where all conditional distributions are from the exponential families. This assumption is required in all our proofs in order to show that an agent's loss exponent cannot be better than a certain achievable lower bound, which we characterize in order to derive the agent's optimal policy.

Assumption 2. For all $m \in [0, M-1]$ and all $\gamma \in \bigcup_k \Gamma_k$, $\xi_m(\gamma, t) < \infty$ for all $t \in \mathbb{R}^{M-1}$.

E. Results Overview

If agent 0 incorporates the opinion of expert k , we expect the loss exponent of expert k to determine how useful its opinion is to agent 0. Therefore, we first find the loss exponent of each expert k in Theorem 2, which also provides an asymptotically optimal policy for expert k . Then, we proceed to determine the loss exponent of agent 0 if it incorporates the opinion of a

particular expert k in Theorem 4. To find the optimal expert, we simply optimize the loss exponent found in Theorem 4 over the whole set of experts, which is essentially the content of Theorem 5. Based on the conclusions of Theorem 5, we provide detailed procedures in Remark 2 of Section IV-A that allow agent 0 to compute its asymptotically optimal expert choice and policy. Finally, we provide insights into the optimal expert choice under simplifying assumptions in Propositions 6 and 7.

III. OPTIMAL LOSS EXPONENTS

In this section, we first characterize the loss exponents of the experts $k = 1, 2, \dots, K$, which leads to an asymptotically optimal policy for each expert. We then characterize the loss exponent of agent 0, assuming that it is following the opinion of some expert k .

A. Loss Exponents of Experts

Consider an expert k , where $k = 1, \dots, K$. By conditioning on the observations $Y_k[1 : n_k]$, it can be shown (Proposition 2.3 of [24]) that the optimal decision rule for expert k is given by

$$D_k = \arg \min_{0 \leq d \leq d_k - 1} \sum_{m=0}^{M-1} \pi_m C_k(m, d, n_k) \ell_{m0}(Y_k[1 : n_k]), \quad (7)$$

where $\ell_{m0}(Y_k[1 : n_k]) = \prod_{l=1}^{n_k} \ell_{m0}(Y_k[l])$. In general, the right hand side of (7) may have multiple minimizers. In order to avoid having to consider the use of randomization to determine the final decision for agent k (due to a mathematical technicality in the proof of Theorem 2 below), and for the ease of interpreting results, we will assume throughout this paper that the right hand side of (7) has a unique solution with probability one.

Assumption 3. *For every agent $k = 1, \dots, K$, and for any policy x_k , the minimization on the right hand side of (7) has a unique solution with probability one.*

Assumption 3 is satisfied if conditioned on any hypothesis $H \in \{0, \dots, M-1\}$, $(\ell_{m0}^\gamma)_{m \geq 1}$ are continuous random variables with a joint probability density function. To see this, it can be shown via an easy inductive argument that $(\ell_{m0}(Y_k[1 : n_k]))_{m \geq 1}$ are in turn continuous random variables under any hypothesis H . Assumption 3 then follows by the same argument in Lemma 5.1 of [13], which we refer the reader to. This shows that Assumption 3 holds in most practical applications as private observations are usually modeled as noisy, with the noise typically having a probability density function like the Gaussian probability density [36], [37].

If agent 0 follows expert k , its loss exponent depends on the loss exponent (2) of the expert k , which is in turn related to the probability exponents

$$\limsup_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{P}_m(D_k = d), \quad (8)$$

where $m \in [0, M-1]$ and $d \in [0, d_k-1]$. In the following, our aim is to characterize the probability exponents (8). We use (7) to characterize (8) by first letting

$$\tilde{g}_k(z, d, n_k) = \frac{1}{n_k} \log \sum_{m=0}^{M-1} \pi_m e^{\log C_k(m, d, n_k) + n_k z[m]},$$

and

$$g_k(z, d, n_k) = \tilde{g}_k(z, d, n_k) - \min_{d' \neq d} \tilde{g}_k(z, d', n_k),$$

where $z = (z[0], z[1], \dots, z[M-1]) \in \mathbb{R}^M$. Suppose that agent k adopts the policy x_{k, n_k} when it has access to n_k private observations. Then from (7) and Assumption 3, we obtain

$$\mathbb{P}_m(D_k = d) = \mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k, n_k}), d, n_k) < 0), \quad (9)$$

where

$$\bar{Z}_{n_k}(x_{k, n_k}) = \left(\frac{1}{n_k} \log \ell_{m0}(Y_k[1 : n_k]) \right)_{m=1}^{M-1} \quad (10)$$

is a vector in \mathbb{R}^M of log likelihood ratios, and $\bar{Z}_{n_k}^0(x_{k, n_k}) = (0, \bar{Z}_{n_k}(x_{k, n_k}))$ is the vector of log likelihood ratios augmented with 0 as the first entry (this corresponds to $1/n_k \cdot \log \ell_{00}(Y_k[1 : n_k]) = 0$). In the following, we show that $g_k(z, d, n_k)$ converges uniformly in z to

$$f_k(z, d) = \tilde{f}_k(z, d) - \min_{d' \neq d} \tilde{f}_k(z, d'), \quad (11)$$

where

$$\tilde{f}_k(z, d) = \max_{0 \leq m \leq M-1} \{z[m] - c_k(m, d)\}. \quad (12)$$

Lemma 1. *Suppose that Assumption 1 holds. For all $k = 1, \dots, K$, and all $d \in [0, d_k-1]$, $\tilde{g}_k(z, d, n) \rightarrow \tilde{f}_k(z, d)$ and $g_k(z, d, n) \rightarrow f_k(z, d)$ uniformly in $z \in \mathbb{R}^M$ as $n \rightarrow \infty$.*

Proof: See Appendix B. ■

We are now ready to present our first main result. Since we are working in the regime of large n_k , Lemma 1 tells us that we can replace g_k with f_k in (9). For $k \geq 1$, let

$$A_k(d) = \{z \in \mathbb{R}^{M-1} : f_k(z^0, d) < 0, \text{ where } z^0 = (0, z)\}, \quad (13)$$

The set $A_k(d)$ can be interpreted as the asymptotic decision region for expert k to declare decision d based on its sufficient statistics $\bar{Z}_{n_k}(x_{k, n_k})$; see Section III-B for a discussion. Assumption 3 is required here to ensure that the sets $A_k(d), d \in [0, d_k-1]$ are disjoint.

Theorem 2. *Suppose that Assumptions 1, 2, and 3 hold. Consider an agent $k \in \{1, \dots, K\}$ who adopts a sequence of optimal policies $(x_{k, n_k})_{n_k \geq 1}$ that minimizes (1) for each n_k . Then, we have the following:*

(i) *The loss exponent of agent k is given by*

$$\lim_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, n_k)] = - \max_{x \in \mathbb{S}(\Gamma_k)} I_k(x),$$

where

$$I_k(x) = \min_{\substack{0 \leq m \leq M-1 \\ 0 \leq d \leq d_k-1}} \left\{ \inf_{z \in A_k(d)} \Phi_m^*(z, x) + c_k(m, d) \right\}. \quad (14)$$

- (ii) There is no loss in optimality asymptotically, if we restrict the sequence of policies $(x_{k,n_k})_{n_k \geq 1}$ such that $\lim_{n_k \rightarrow \infty} x_{k,n_k} = x_k^*$, for some $x_k^* \in \arg \max_{x \in \mathbb{S}(\Gamma_k)} I_k(x)$.
- (iii) Let $x_k^* \in \arg \max_{x \in \mathbb{S}(\Gamma_k)} I_k(x)$. For every $m \in [0, M - 1]$, and $d \in [0, d_k - 1]$, we have

$$\lim_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) = - \inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*).$$

Proof: See Appendix B. \blacksquare

Theorem 2 allows each expert to find an asymptotically optimal policy by maximizing (14) (see Remark 2 in Section IV-A). In addition, we will see later that Theorem 2(iii) allows agent 0 to characterize its own loss exponent and thus optimally choose the expert to follow.

Remark 1. The quantity $I_k(x)$ in (14) can be interpreted as the rate of loss decay for agent k adopting policy x , i.e., for n_k sufficiently large, we have

$$\mathbb{E}[C_k(H, D_k, n_k)] \approx g(n_k) e^{-I_k(x)},$$

where $g(n_k)$ is a function that decays faster than exponentially in n_k . To achieve a smaller loss, agent k then chooses a policy x that maximizes $I_k(x)$. Observe that $I_k(x)$ is the minimization over all $m \in [0, M - 1]$ and $d \in [0, d_k - 1]$ of the terms $\inf_{z \in A_k(d)} \Phi_m^*(z, x) + c_k(m, d)$, where $\inf_{z \in A_k(d)} \Phi_m^*(z, x)$ is the absolute probability exponent in Theorem 2(iii), while $c_k(m, d)$ is the loss function decay rate of declaring decision d when the true hypothesis is m . Therefore, each term in the minimization of the right hand side of (14) is the absolute decay rate of $C_k(m, d, n_k) \mathbb{P}_m(D_k = d)$, and the overall loss exponent is dominated by the term that decays the slowest, which is what we expect from large deviations theory (cf. Lemma 1.2.15 of [35]). We next provide an intuitive interpretation for the regions $A_k(d)$ in the following subsection.

B. Asymptotic Decision Regions

Suppose that agent $k \geq 1$ adopts the sequence of policies $x_{k,n_k} \rightarrow x_k$ as $n_k \rightarrow \infty$. From (9), for large n_k , we have for each $m \in [0, M - 1]$, and $d \in [0, d_k - 1]$,

$$\begin{aligned} \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) &\approx \frac{1}{n_k} \log \mathbb{P}_m(f_k(\bar{Z}_{n_k}(x_{k,n_k}), d) < 0) \\ &\approx - \inf_{z \in A_k(d)} \Phi_m^*(z, x_k), \end{aligned}$$

where $A_k(d)$ is as defined in (13), and $\Phi_m^*(z, x_k)$ can be interpreted as a rate function. In the same spirit as the Gärtner-Ellis Theorem [35], we can interpret the sets $A_k(d)$ as the asymptotic decision region for deciding $D_k = d$ based on $\bar{Z}_{n_k}(x_{k,n_k})$ as $n_k \rightarrow \infty$, where the rate of decay of $\mathbb{P}_m(D_k = d)$ is dominated by the rate at a particular realization of $\bar{Z}_{n_k}(x_{k,n_k})$ in the region $A_k(d)$. We refer the reader to Appendix C for a characterization of $A_k(d)$ in terms of the union of intersections of multiple half-spaces.

In the following, we list some properties of the rate functions $\Phi_m^*(\cdot, \cdot)$ that will be useful in helping us to interpret our results.

Lemma 3. For every $m \in [0, M - 1]$, $\Phi_m^*(z, x)$ is non-negative, convex in z and concave in x . Furthermore, for any policy x , $\min_{z \in \mathbb{R}^{M-1}} \Phi_m^*(z, x) = 0$, and the minimum is achieved at

$$\tilde{z}_m(x) = \sum_{\gamma} x[\gamma] \mathbb{E}_m[Z^\gamma], \quad (15)$$

where $Z^\gamma = (\log \ell_{i0}^\gamma)_{i=1}^{M-1}$.

Proof: See Appendix B. \blacksquare

In particular, if $d_k = M$, and the loss function of agent k adopting the policy x_k is such that $\tilde{z}_m(x_k) \in A_k(m)$ in (15) for every m , then we can interpret the decision m of agent k as in favor of hypothesis $H = m$ since $\mathbb{P}_m(D_k = m)$ is bounded away from 0 and has decay rate $\inf_{z \in A_k(m)} \Phi_m^*(z, x_k) = 0$.

C. Loss Exponent of Agent 0

In the following, we characterize the loss exponent of agent 0 if it chooses expert k . Recall that $q_k = \lim_{n_0 \rightarrow \infty} n_k/n_0$.

Theorem 4. Suppose that Assumptions 1, 2, and 3 hold. Suppose that agent 0 adopts the opinion of expert $k \geq 1$, which has the asymptotically optimal policy x_k^* . Then, the loss exponent of agent 0 is

$$\lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{E}[C_0(H, D_0(k), n_0)] = - \max_{x_0 \in \mathbb{S}(\Gamma_0)} \mathcal{E}_0(k, x_0),$$

where ⁵

$$\begin{aligned} \mathcal{E}_0(k, x_0) &= \min_{\substack{i \neq j \\ 0 \leq d \leq d_k - 1}} \max_{s \in [0, 1]} \left\{ (1-s) \left(q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) + c_0(i) \right) \right. \\ &\quad \left. + s \left(q_k \inf_{z \in A_k(d)} \Phi_j^*(z, x_k^*) + c_0(j) \right) - \Lambda_{ij}(s, x_0) \right\}. \end{aligned} \quad (16)$$

Proof: See Appendix B. \blacksquare

Following Remark 1, we can again interpret each term in the minimization on the right hand side of (16) as the absolute error exponent incurred when differentiating between hypotheses i and j . From Theorem 2 and Theorem 4, we see that the choice of an expert k affects agent 0's loss exponent through the rate of decay of the probabilities $\mathbb{P}_i(D_k = d)$ and $\mathbb{P}_j(D_k = d)$, which is to be expected. In particular, if the loss function of the expert k is such that we can interpret the decision d of the expert as in favor of hypothesis d , then $\mathbb{P}_i(D_k = i) > 0$ yields $\inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) = 0$, and it does not contribute to the decay of the expected loss of agent 0.

IV. OPTIMAL CHOICE OF EXPERT

We now address the question of how to choose an optimal expert for agent 0. We first revisit Theorem 4 to derive the optimal loss exponent for agent 0, and then we make additional simplifying assumptions in order to provide more insights into our results. Finally, we discuss a numerical example to illustrate the use of our optimal loss exponent characterizations in finding the optimal policies and expert choice.

⁵To avoid cluttered notations, we let $\min_{i \neq j}$ and $\max_{i \neq j}$ be the minimization or maximization over all unordered pairs $(i, j) \in [0, M - 1]^2$ such that $i \neq j$, respectively.

A. Choosing an Expert

If agent 0 does not incorporate the opinion of any expert, we use the notation $\mathcal{E}_0(0, x_0)$ to denote its absolute loss exponent when using policy x_0 . From Theorem 4, by setting $q_k = 0$, we have

$$\mathcal{E}_0(0, x_0) = \min_{i \neq j} \max_{s \in [0,1]} \{(1-s)c_0(i) + sc_0(j) - \Lambda_{ij}(s, x_0)\}. \quad (17)$$

In the case of minimizing the error probability at agent 0 with 0-1 loss function, and without the help of any experts, we can further set $c_0(i) = c_0(j) = 0$ to obtain the absolute error exponent

$$\mathcal{E}_{0,B}(0, x_0) = -\max_{i \neq j} \min_{s \in [0,1]} \Lambda_{ij}(s, x_0), \quad (18)$$

which recovers the result in [32]. Similarly, if agent 0 chooses expert k , and adopts the 0-1 loss function, we let its absolute loss exponent be

$$\begin{aligned} \mathcal{E}_{0,B}(k, x_0) = & \min_{\substack{i \neq j \\ 0 \leq d \leq d_k - 1}} \max_{s \in [0,1]} \left\{ (1-s)q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) \right. \\ & \left. + sq_k \inf_{z \in A_k(d)} \Phi_j^*(z, x_k^*) - \Lambda_{ij}(s, x_0) \right\}. \end{aligned} \quad (19)$$

The following proposition follows immediately from Theorem 4, and provides a method for agent 0 to optimally choose which expert to follow.

Theorem 5. Suppose that Assumptions 1, 2, and 3 hold. The optimal loss exponent of agent 0 is

$$\begin{aligned} & \min_{1 \leq k \leq K} \lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{E}[C_0(H, D_0(k), n_0)] \\ & = -\max_{1 \leq k \leq K} \max_{x_0 \in \mathbb{S}(\Gamma_0)} \mathcal{E}_0(k, x_0). \end{aligned}$$

Furthermore, for any $k \geq 1$ and policy x_0 , we have

$$\mathcal{E}_0(k, x_0) \geq \mathcal{E}_{0,B}(k, x_0) \quad (20)$$

$$\mathcal{E}_0(k, x_0) \geq \mathcal{E}_0(0, x_0) \geq \mathcal{E}_{0,B}(0, x_0). \quad (21)$$

Proof: See Appendix B. ■

Remark 2. Theorem 5 shows that agent 0 should choose an expert k that maximizes $\max_{x_0} \mathcal{E}_0(k, x_0)$, which depends on $\inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*)$, for $m \in [0, M-1]$ and $d \in [0, d_k-1]$. We now discuss procedures that can achieve this, depending on the amount of information that agent 0 has about the experts:

- (i) If the experts' probability exponents $\{\inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*) : m \in [0, M-1], d \in [0, d_k-1]\}$ for all $k \geq 1$ are publicized, then agent 0 simply needs to find its optimal policy x_0 that maximizes $\mathcal{E}_0(k, x_0)$ in Theorem 5 for each k , and choose the expert that produces the largest $\max_{x_0} \mathcal{E}_0(k, x_0)$. To find the optimal policy x_0 corresponding to an expert k , we can use an iterative procedure as follows:

Step 1. An initial guess for the optimal policy $x_0^*(0)$ is made. Set $l = 1$.

Step 2. For each pair $(i, j) \in [0, M-1]$, $i \neq j$, and $d \in [0, d_k-1]$, find $s_{ijd}(l)$ that is the maximizer

of the optimization problem over $s \in [0, 1]$ on the right hand side of (16). Note that this is equivalent to a convex minimization problem, which can be solved via standard convex optimization methods [38].

Step 3. Find $x_0^*(l) = (x[\gamma])_{\gamma \in \Gamma_0}$ by solving the linear program:

$$\begin{aligned} & \max_{r, (x[\gamma])_{\gamma \in \Gamma_0}} r \\ & \text{subject to} \\ & r \leq (1 - s_{ijd}(l)) \left(q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) + c_0(i) \right) \\ & \quad + s_{ijd}(l) \left(q_k \inf_{z \in A_k(d)} \Phi_j^*(z, x_k^*) + c_0(j) \right) \\ & \quad - \sum_{\gamma \in \Gamma_0} x[\gamma] \log \mathbb{E}_i[(\ell_{ji}^\gamma)^{s_{ijd}(l)}], \\ & \quad \forall i, j \in [0, M-1], i \neq j, \forall d \in [0, d_k-1], \\ & \quad x[\gamma] \geq 0, \forall \gamma \in \Gamma_0, \\ & \quad \sum_{\gamma \in \Gamma_0} x[\gamma] = 1. \end{aligned}$$

Step 4. Set $l = l + 1$ and repeat Steps 2-4 till $x_0^*(l)$ does not change significantly.

- (ii) If agent 0 knows only the experts' loss functions and optimal policies x_k^* for all $k \geq 1$, then it can compute $\inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*)$ for all $m \in [0, M-1]$ and $d \in [0, d_k-1]$ by searching for the minimum of $\Phi_m^*(z, x_k^*)$ on the boundaries of $A_k(d)$, if $\tilde{z}_m(x_k^*)$ in (15) is not in $A_k(d)$. This is because $\Phi_m^*(z, x_k^*)$ is convex in z (cf. Lemma 3), and the search can be performed using standard convex optimization methods [38]. The boundaries of $A_k(d)$ can be found through its characterization in Appendix C (see Section IV-C for an example). If $\tilde{z}_m(x_k^*) \in A_k(d)$, then from Lemma 3, we obtain $\inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*) = 0$. Once the experts' probability exponents have been computed, the same procedure as in item (i) above can now be used to choose the optimal expert.

- (iii) If agent 0 knows only the experts' loss functions, it can utilize Theorem 2(i) to determine the policy that each expert adopts. This however may not be an easy numerical procedure if M and d_k are large, even in the case where the expert has 0-1 loss function [32]. We propose the following alternating optimization approach for finding the optimal policy of expert k :

Step 1. An initial guess for the optimal policy $x_k^*(0)$ is made. Set $l = 1$.

Step 2. For each $m \in [0, M-1]$ and $d \in [0, d_k-1]$, find $z_{m,d}(l) = \arg \inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*(l-1))$. As noted in item (ii) above, this can be obtained via standard convex optimization methods. Let

$$\begin{aligned} t_{m,d}(l) = & \arg \max_{t \in \mathbb{R}^{M-1}} \{ \langle t, z_{m,d}(l-1) \rangle \\ & - \varphi_m(t, x_k^*(l)) \}. \end{aligned}$$

Step 3. Find $x_k^*(l) = (x[\gamma])_{\gamma \in \Gamma_k}$ by solving the linear program

$$\begin{aligned} & \max_{r, (x[\gamma])_{\gamma \in \Gamma_k}} r \\ & \text{subject to} \\ & r \leq \langle t_{m,d}(l), z_{m,d}(l) \rangle \\ & \quad - \sum_{\gamma \in \Gamma_k} x[\gamma] \xi_m(\gamma, t_{m,d}(l)) + c_k(m, d), \\ & \quad \forall m \in [0, M-1], \forall d \in [0, d_k-1], \\ & x[\gamma] \geq 0, \forall \gamma \in \Gamma_k, \\ & \sum_{\gamma \in \Gamma_k} x[\gamma] = 1. \end{aligned}$$

Step 4. Set $l = l + 1$ and repeat Steps 2-4 till $x_k^*(l)$ does not change significantly.

Unfortunately, there is no guarantee that the above procedure converges to the correct solution. However, in our numerical experiments, we were able to arrive at the correct optimal policy by using multiple initial guesses. A numerical example is presented in Section IV-C.

Remark 3. The first inequality in (21) shows that there is no loss in optimality for agent 0 to incorporate the opinion of any expert, verifying the adage that there is no harm in having more information. However, incorporating additional information does not necessarily improve its loss exponent; see Proposition 6.

The inequality in (20) shows that of all the loss functions satisfying Assumption 1(i), the 0-1 loss is the worst canonical loss function for agent 0. In particular, consider the set of loss functions with $C_0(m, d, n_0) = C_0(m, n_0)$ for all $m \in [0, M-1]$ and $d \neq m$, where the limit $c_0(m) = -\lim_{n_0 \rightarrow \infty} (1/n_0) \log C_0(m, n_0)$ exists. Furthermore, each loss function has a total cost constraint,

$$\sum_{m=0}^{M-1} C_0(m, n_0) = M - 1.$$

It can be shown that all such loss functions satisfy Assumption 1(i), and are canonical loss functions. The 0-1 loss function belongs to this set, and divides the total cost equally among all types of missed detections, which results in the worst loss exponent for agent 0. This can be explained intuitively by observing that if there exists a m such that $c_0(m) > 0$, then missed detection of $H = m$ incurs an exponentially decaying loss, so that the hypothesis m can effectively be ignored. Agent 0 then effectively has a smaller set of hypotheses, leading to a lower expected loss.

Remark 4. It is easy to generalize Theorem 4 or Theorem 5 to the case where more than one agent can be chosen. In particular, if all K experts' opinions are adopted by agent 0, Theorem 4 holds with $\mathcal{E}_0(k, x_0)$ replaced by

$$\begin{aligned} & \min_{(p_k)_{k=1}^K} \max_{i \neq j, s \in [0,1]} \left\{ (1-s) \left(\sum_{k=1}^K q_k \inf_{z \in A_k(p_k)} \Phi_i^*(z, x_k^*) + c_0(i) \right) \right. \\ & \quad \left. + s \left(\sum_{k=1}^K q_k \inf_{z \in A_k(p_k)} \Phi_j^*(z, x_k^*) + c_0(j) \right) - \Lambda_{ij}(s, x_0) \right\} \end{aligned}$$

where the sequences $(p_k)_{k=1}^K \in \prod_{k=1}^K [0, d_k - 1]$.

In the following subsection, we consider the special cases where the size of each expert's decision space is less than or equal to M under additional simplifying assumptions.

B. Special Cases

In this section, we assume that agent 0's policy has been fixed in advance. We first consider the case where an expert k may have a decision space smaller than the number of hypotheses M that agent 0 has. To state our results, suppose that agent 0 adopts policy x_0 and has loss decay rates given by the function c_0 . We define the *loss augmented* log moment generating function for $i, j \in [0, M-1]$ as

$$\bar{\Lambda}_{ij}(s, x_0, c_0) = \Lambda_{ij}(s, x_0) - (1-s)c_0(i) - sc_0(j). \quad (22)$$

For the case where agent 0 adopts the 0-1 loss function, we have $c_0(i) = 0$ for all i . We observe that in this case, the Chernoff information $-\min_{s \in [0,1]} \bar{\Lambda}_{ij}(s, x_0)$ tells us how well hypothesis i can be differentiated from hypothesis j , where a larger Chernoff information corresponds to a faster error probability decay [35]. This leads us to the following general definition: For a given policy x_0 , if there exists a constant α so that $\min_{s \in [0,1]} \bar{\Lambda}_{ij}(s, x_0, c_0) = \alpha$ for all $i \neq j$, we say that the policy x_0 is *hypothesis-loss neutral* for agent 0.

Intuitively, a policy is hypothesis-loss neutral if it can differentiate any pair of hypotheses equally well when adjusted for the agent's "biases". It may be argued that in practical scenarios, an agent would strive to be hypothesis-loss neutral, since otherwise there are hypotheses that are relatively less important than others, and can be dropped. In the following, we show that if agent 0 is hypothesis-loss neutral, an expert's opinion is useless if it does not have a decision space as large as M .

Proposition 6. *Suppose that Assumptions 1, 2, and 3 hold. Suppose that agent 0 adopts a policy x_0 that is hypothesis-loss neutral. Then, for any agent $k \geq 1$, if $d_k < M$, we have $\mathcal{E}_0(k, x_0) = \mathcal{E}_{0,B}(0, x_0)$, i.e., agent 0 ignores the opinion of agent k .*

Proof: See Appendix B. ■

Proposition 6 can be explained intuitively as follows in the case where agent 0 has 0-1 loss function. If an expert k has $d_k < M$, then there exists two hypotheses i and j of agent 0 that it does not discriminate between. The probability of making an error between these two hypotheses by the expert is therefore one, and the expert's opinion does not help agent 0 to differentiate between hypotheses i and j . The error probability of agent 0 declaring $H = i$ when the true hypothesis is $H = j$ or vice versa, thus dominates when agent 0 is hypothesis-loss neutral, and is the same as when agent 0 ignores the expert's opinion.

We next consider the case where experts have the same decision space as agent 0, and has zero loss if they decide on the true underlying hypothesis. We also make the following simplifying assumption.

Assumption 4. Every agent $k \geq 1$ has $d_k = M$, with loss decay rates

$$c_k(m, d) = \begin{cases} c_k(m) & \text{if } d \neq m, \\ \infty & \text{if } d = m, \end{cases}$$

where $c_k(m) \geq 0$ for all $m \in [0, M - 1]$.

Proposition 7. Suppose that Assumptions 1-4 hold. Suppose that agent 0 adopts policy x_0 , and follows the opinion of expert $k \geq 1$, who adopts the policy x_k . Then, the loss exponent of agent 0 is

$$\lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{E} [C_0(H, D_0(k), n_0)] = -\tilde{\mathcal{E}}_0(k, x_0), \quad (23)$$

where ⁶

$$\tilde{\mathcal{E}}_0(k, x_0) = \min_{i, j: i \neq j} \max_{s \in [0, 1]} \{sq_k \Lambda_{ji}^*(c_k(i) - c_k(j), x_k) - \bar{\Lambda}_{ij}(s, x_0, c_0)\}.$$

In addition, if for some $i \neq j$ such that $c_0(i) = c_0(j)$, and

$$\Lambda_{ij}(s, x_0) = \Lambda_{ji}(s, x_0), \quad \forall s \in [0, 1], \quad (24)$$

then there is no loss in optimality for agent 0 to restrict to experts k with canonical loss functions satisfying $c_k(i) = c_k(j)$.

Proof: See Appendix B. ■

Surprisingly, Proposition 7 shows that it is not necessarily optimal for agent 0 to choose an expert who utilizes the same canonical loss function as itself since agent 0's loss exponent depends only on the relative differences in agent k 's loss function at each hypothesis. We show a numerical example of this phenomenon in Section IV-C below. In particular, if agent 0 adopts the 0-1 loss function, and (24) does not hold for some $i \neq j$, then an example can be constructed in which an expert with a loss function different from the 0-1 loss is optimal for agent 0.

On the other hand, if agent 0 is “unbiased” (in terms of loss) towards a pair of hypotheses (i, j) and (24) holds, Proposition 7 tells us that it is optimal for agent 0 to choose an expert, if any, who is also “unbiased” towards hypotheses (i, j) . This indicates that if agent 0 has the same discriminatory power for a particular pair of hypotheses (which is the intuitive meaning of (24)), then it values information from an “unbiased” expert more than one who has the same “bias” as itself. Assuming that independent news agencies are in general “unbiased,” our result suggests that such news agencies are unlikely to be replaced by social news reporting or social blogs. The following corollary follows immediately from Proposition 7.

Corollary 8. Suppose that Assumptions 1-4 hold. Suppose that agent 0 adopts policy x_0 , and has 0-1 loss function. If for every distinct pair of hypotheses (i, j) , (24) holds, then it is optimal for agent 0 to choose an expert, if any, who also has the 0-1 loss function.

⁶The notation $\min_{i, j: i \neq j}$ means minimization over all ordered pairs (i, j) with $i, j \in [0, M - 1]$ and $i \neq j$.

C. Numerical Example

In this section, we present a numerical example to illustrate our results. Suppose that $M = 3$, there are 3 experts to choose from, and all agents have access to private observations from two information sources $\Gamma = \{\gamma_1, \gamma_2\}$, whose conditional distributions are shown in Table I. We use the notation $\mathcal{N}(\mu, \sigma^2)$ to denote the normal distribution with mean μ and variance σ^2 .⁷ For example, the three hypotheses can correspond to the financial health of a company being “bad”, “neutral”, and “good” respectively. The information source γ_1 represents a view that tends to be optimistic when the true health of the company is “neutral”, while γ_2 represents a view that tends to be pessimistic.

TABLE I: Conditional distributions of the information sources.

Hypothesis H	γ_1	γ_2
0	$\mathcal{N}(-1, \sigma^2)$	$\mathcal{N}(-1, \sigma^2)$
1	$\mathcal{N}(\delta, \sigma^2)$	$\mathcal{N}(-\delta, \sigma^2)$
2	$\mathcal{N}(1, \sigma^2)$	$\mathcal{N}(1, \sigma^2)$

In this example, we take $\sigma = 2$, $q_k = 1$ for all $k \geq 1$, and assume Assumption 4 for easier interpretation of our results. The asymptotic decision regions of each agent $k \geq 1$ can be shown, using the characterization in Appendix C, to be that in Figure 2. For each agent k , let $c_k = (c_k(0), c_k(1), c_k(2))$ be its vector of loss decay rates. Let $c_1 = (0, 0, 0)$, $c_2 = (0, 0, 0.2)$, and $c_3 = (0, 0.05, 0)$. We use the alternating optimization procedure in Remark 2 item (iii) to find the optimal policies for each agent. An exhaustive search is done to determine the actual optimal policy for each expert, which is given in Table II. Note that the experts' optimal policies are invariant of δ .




FIG. 2: Asymptotic decision regions for agent k . The dotted lines denote the boundaries between the regions.

For experts 1 and 3, their optimal policies are both given by $x_1[\gamma_1] = x_1[\gamma_2] = 0.5$, as is expected from the symmetry of the problem. For expert 2, its optimal policy is given by $x_2[\gamma_1] = 1$ and $x_2[\gamma_2] = 0$. This is because with a positive loss decay rate $c_2(2) = 0.2$, expert 2 puts exponentially less loss on confusing hypothesis 2 for the others, therefore it chooses γ_1 as its sole information source. For the alternating optimization

⁷Normal distributions have been used to model social opinions and influences in the economics literature, including [36], [37].

TABLE II: Optimal policies for agents. The leftmost columns show the optimal policies for each expert. The three rightmost columns show the optimal policy and loss decay rate of an agent 0 with $c_0 = c_3$, corresponding to each expert choice.

Expert k	$x_k^*[\gamma_1]$	$x_k^*[\gamma_2]$	$x_0^*[\gamma_1]$	$x_0^*[\gamma_2]$	$\mathcal{E}_0(k, x_0)$
1	0.5	0.5	0.5	0.5	0.1099
2	1	0	0.2117	0.7883	0.1158
3	0.5	0.5	0.5	0.5	0.1066

procedure, we started our initial guesses at $(x[\gamma_1], x[\gamma_2]) = (0, 1)$ and $(0.3, 0.7)$.⁸ We repeated the procedure for different values of δ . In all cases, we were able to find the correct optimal policies within 22 iterations (totaled over the two runs corresponding to the two initial guesses), as shown in Figure 3. In some cases however, the procedure does not converge to the correct optimal policy if we restrict ourselves to the single initial guess $(x[\gamma_1], x[\gamma_2]) = (0, 1)$. This shows that the proposed procedure can get stuck at a suboptimal policy.




FIG. 3: Total number of iterations required to find optimal policies for experts 1, 2, and 3.

We next suppose that $c_0 = c_3$, and use the procedure in Remark 2 item (i) to obtain the optimal policy for agent 0 w.r.t. each expert $k = 1, 2, 3$, when $\delta = 0.9$. We found the optimal policies for agent 0 within 4 iterations in each case, and that choosing expert 2 produces the best loss exponent for agent 0, as shown in Table II. In this case, it is not optimal for agent 0 to choose expert 3, which has the same loss function as itself. Observe that even if agent 0's policy is fixed at $x_0 = (0.5, 0.5)$, it will still prefer expert 1 over expert 3. On the other hand, if $c_0 = c_1$, we have $\mathcal{E}_0(1, x_0) = 0.0884 > \mathcal{E}_0(3, x_0) = 0.0750 > \mathcal{E}_0(2, x_0) = 0.0566$, i.e., agent 0's optimal expert choice is expert 1. It can be shown that (24) holds for all pairs of hypotheses in this example, thus verifying Corollary 8.

V. CONCLUSION

We have studied a multihypothesis social learning problem in which an agent makes a decision with the help of a chosen expert. We have considered a general framework that allows the agent and experts to have different loss functions (biases), and different decision spaces. We have characterized the loss exponent of the agent in terms of the chosen expert's probability error exponent, which allows us to choose the asymptotically optimal expert as well as the agent's policy.

⁸Although a more natural initial guess is $(0.5, 0.5)$, this was not selected in order to test the convergence for experts 1 and 3.

We have shown that if the experts have the same decision space as the agent, then it is not necessarily optimal for the agent to choose an expert with the same loss function as itself. Moreover, if the agent is hypothesis-loss neutral, then it ignores any expert with a decision space smaller than the number of hypotheses.

The results in this paper are limited by our setup and assumptions, which however allows us to obtain analytical characterizations of the asymptotically optimal policies for the agent and experts, and an asymptotically optimal expert choice. In the regime of a finite number of private observations, these policies are in general sub-optimal. Finding the exact optimal policies in this regime is however analytically difficult, and the asymptotically optimal policies can be utilized when the numbers of private observations are large.

In this paper, we have assumed that the agent knows the experts' loss decay rates, which may not be valid in some practical scenarios. For example, expert opinions may very well depend on the mood of the expert at the time the opinion is publicized. Therefore, it is of interest to consider minimax loss exponents in which the expert's loss function or observation probability distributions are drawn from uncertainty classes. Minimax decentralized hypothesis testing has been studied in [39], and in our recent work [34] in which we consider robust social learning in a tandem network. However, additional research in minimax decentralized hypothesis testing is required to address more complex network architectures and applications like that considered in this paper.

APPENDIX A MATHEMATICAL PRELIMINARIES

In this appendix, we briefly review some basic definitions and a result from large deviations theory. The reader is referred to [35], [40] for details. The notation $\text{Tr}(A)$ denotes the trace of the matrix A , $\nabla_t f(t)$ is the gradient of $f(t)$ w.r.t. the vector t , and $\nabla_t^2 f(t)$ is the Hessian w.r.t. t of the function $f(t)$.

Let \mathcal{X} be a Polish space. The function $\Phi : \mathcal{X} \mapsto [0, \infty]$ is said to be a good rate function if Φ is lower semicontinuous and has compact level sets $\{x \in \mathcal{X} : \Phi(x) \leq c\}$ for all $c \geq 0$. Let $(\mathbb{P}_n)_{n \geq 1}$ be a sequence of probability measures on \mathcal{X} . This sequence of probability measures is said to satisfy a *large deviation principle* (LDP) if for some good rate function Φ , and for all closed sets $C \subset \mathcal{X}$, we have

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_n(C) \leq - \inf_{x \in C} \Phi(x),$$

and for all open sets $O \subset \mathcal{X}$, we have

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_n(O) \geq - \inf_{x \in O} \Phi(x).$$

A sequence of random variables $(Z_n)_{n \geq 1}$ is said to satisfy a LDP if the sequence of marginal distributions $\mathbb{P}_n(\cdot) = \mathbb{P}(Z_n \in \cdot)$ satisfies a LDP. The celebrated Gärtner-Ellis Theorem [35] provides sufficient conditions for a sequence of random variables to satisfy a LDP. In the following, we present a partial version of the Gärtner-Ellis Theorem

Theorem A.1. Let $(Z_n)_{n \geq 1}$ be a sequence of \mathbb{R}^M -valued random variables, so that

$$\varphi(t) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}[\exp(\langle t, Z_n \rangle)]$$

exists. Let

$$\Phi^*(z) = \sup_t \{ \langle t, z \rangle - \varphi(t) \}$$

be the Fenchel-Legendre transform of φ . Then, under the assumptions in Theorem 2.3.6 of [35], the LDP holds with good rate function Φ^* .

When not all the conditions required by the Gärtner-Ellis Theorem hold (as is the case in some of our proofs), we instead make use of a *uniform* lower bound given by the following lemma, which is a generalization of Theorem 1.3.13 of [40]. The proof is omitted here for brevity.

Lemma A.1. Let Z_1, Z_2, \dots, Z_n be independent \mathbb{R}^M -valued random variables, with $x[\gamma]$ fraction of them having distribution \mathbb{P}^γ , for each $\gamma \in \Gamma$. Let Z^γ have distribution \mathbb{P}^γ ,

$$\varphi(t) = \sum_{\gamma \in \Gamma} x[\gamma] \log \mathbb{E}[\exp(\langle t, Z^\gamma \rangle)],$$

and

$$\Phi^*(z) = \sup_t \{ \langle t, z \rangle - \varphi(t) \}.$$

Suppose that $\mathbb{E}[\exp(\langle t, Z^\gamma \rangle)] < \infty$ for all $t \in \mathbb{R}^M$, then for t_z such that $\nabla_t \varphi(t_z) = z$, and any $\epsilon > 0$, we have

$$\begin{aligned} & \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} \sum_{l=1}^n Z_l \in B_\epsilon(z) \right) \\ & \geq -\Phi^*(z) - \|t_z\| \epsilon + \frac{1}{n} \log \left(1 - \frac{1}{n\epsilon^2} \text{Tr}(\nabla_t^2 \varphi(t_z)) \right). \end{aligned} \quad (25)$$

In this paper, we apply Lemma A.1 in cases where φ corresponds to φ_m in (4). We will often need to further lower bound (25) by finding an upper bound for $\nabla_t^2 \varphi_m(t_z)$ that does not depend on n or the particular policy adopted by an agent. Assumption 2 implies an upper bound for $\nabla_t^2 \varphi_m$, as shown in the following elementary lemma.

Lemma A.2. Suppose that Assumption 2 holds. Then there exists a non-decreasing function $G(r)$, finite for each $r \geq 0$, such that for all $t \in B_r(0)$, and all $\gamma \in \bigcup_k \Gamma_k$, we have for all $m \geq 0$,

$$0 \leq \text{Tr}(\nabla_t^2 \xi_m(\gamma, t)) \leq G(r). \quad (26)$$

Proof: The lower inequality in (26) holds because $\xi_m(\gamma, t)$ is convex in t for each γ (see Lemma 2.2.31 of [35]). Furthermore, we can define

$$G(r) = \max_{\substack{t \in B_r(0) \\ \gamma \in \Gamma}} \text{Tr}(\nabla_t^2 \xi_m(\gamma, t)),$$

which is finite because of Assumption 2, and the fact that $\text{Tr}(\nabla_t^2 \xi_m(\gamma, t))$ is continuous on the compact set $B_r(0)$. The proof of the lemma is now complete. ■

APPENDIX B PROOFS OF MAIN RESULTS

A. Proof of Lemma 1

Let ϵ be a positive number. From Assumption 1(ii), for n sufficiently large, we have $(\log M)/n \leq \epsilon/2$, $\min_m (\log \pi_m)/n \geq -\epsilon/2$, and

$$\left| \frac{1}{n} \log C_k(m, d, n) + c_k(m, d) \right| \leq \frac{\epsilon}{2},$$

for all $m \in [0, M-1]$. This implies that for all $z \in \mathbb{R}^M$, we have

$$\begin{aligned} \tilde{g}_k(z, d, n) & \leq \max_m \left\{ \frac{1}{n} \log C_k(m, d, n) + z[m] \right\} + \frac{1}{n} \log M \\ & \leq \max_m \{ -c_k(m, d) + z[m] \} + \epsilon \\ & = \tilde{f}_k(z, d) + \epsilon, \end{aligned}$$

and

$$\begin{aligned} \tilde{g}_k(z, d, n) & \geq \max_m \left\{ \frac{1}{n} \log C_k(m, d, n) + z[m] \right\} + \min_m \frac{\log \pi_m}{n} \\ & \geq \max_m \{ -c_k(m, d) + z[m] \} - \epsilon \\ & = \tilde{f}_k(z, d) - \epsilon, \end{aligned}$$

which shows that $\tilde{g}_k(z, d, n) \rightarrow \tilde{f}_k(z, d)$ uniformly in z . Since the minimum of a finite set of uniformly convergent functions is also uniformly convergent, the lemma follows.

B. Proof of Theorem 2

We prove claim (i) of Theorem 2 by first deriving a lower bound for the loss exponent of agent k , and showing that this bound is achievable. We have for any $\epsilon > 0$, and n_k sufficiently large,

$$\begin{aligned} & \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, n_k)] \\ & = \frac{1}{n_k} \log \sum_{m=0}^{M-1} \sum_{d=0}^{d_k-1} \pi_m C_k(m, d, n_k) \mathbb{P}_m(D_k = d) \\ & \geq \max_{\substack{0 \leq m \leq M-1 \\ 0 \leq d \leq d_k-1}} \left\{ \frac{1}{n_k} \log C_k(m, d, n_k) + \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \right\} \\ & \quad + \min_m \frac{1}{n_k} \log \pi_m \\ & \geq \max_{\substack{0 \leq m \leq M-1 \\ 0 \leq d \leq d_k-1}} \left\{ \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) - c_k(m, d) \right\} - \epsilon, \end{aligned} \quad (27)$$

which can be further lower bounded by lower bounds on the probability exponents. For each $z \in \mathbb{R}^{M-1}$, let t_z be the solution to the equation

$$\nabla_t \varphi_i(t, x_{k, n_k}) = z,$$

if the solution exists. For each $r > 0$, let $H_r = \{z \in \mathbb{R}^{M-1} : t_z \text{ exists, } \|t_z\| \leq r\}$, and $A_k(d, r, \epsilon) = \{z \in \mathbb{R}^{M-1} :$

$f_k(z^0, d) < -\epsilon/2$, where $z^0 = (0, z) \cap H_r$. From (9) and (10), we then have for any $u < -\epsilon$,

$$\begin{aligned} & \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \\ &= \frac{1}{n_k} \log \mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d, n_k) < 0) \\ &\geq \frac{1}{n_k} \log \mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d, n_k) \in [u - \epsilon, u + \epsilon]) \\ &\geq \frac{1}{n_k} \log \mathbb{P}_m(f_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d) \in [u - \frac{\epsilon}{2}, u + \frac{\epsilon}{2}]) \end{aligned} \quad (28)$$

where the last inequality follows from Lemma 1 for n_k sufficiently large.⁹ Since $u < -\epsilon$ is arbitrary, we obtain for n_k sufficiently large,

$$\begin{aligned} & \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \\ &\geq \sup_{z: B_\epsilon(z) \subset A_k(d, r, \epsilon)} \frac{1}{n_k} \log \mathbb{P}_m(\bar{Z}_{n_k}(x_{k,n_k}) \in B_\epsilon(z)), \end{aligned} \quad (29)$$

where $B_\epsilon(z)$ is an open sphere of radius ϵ around z . From Lemma A.1, we can further lower bound the right hand side of (29) to obtain

$$\begin{aligned} & \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \\ &\geq - \inf_{z: B_\epsilon(z) \subset A_k(d, r, \epsilon)} \left\{ \Phi_m^*(z, x_{k,n_k}) + \|t_z\| \epsilon \right. \\ &\quad \left. - \frac{1}{n_k} \log \left(1 - \frac{1}{n_k \epsilon^2} \text{Tr}(\nabla_t^2 \varphi_i(t_z, x_{k,n_k})) \right) \right\} \\ &\geq - \inf_{z: B_\epsilon(z) \subset A_k(d, r, \epsilon)} \Phi_m^*(z, x_{k,n_k}) + r \epsilon \\ &\quad - \frac{1}{n_k} \log \left(1 - \frac{1}{n_k \epsilon^2} G(r) \right), \end{aligned} \quad (30)$$

where the last inequality follows from Lemma A.2. Combining (30) with (27), and letting $n_k \rightarrow \infty$, and then taking $r = 1/\sqrt{\epsilon}$ and $\epsilon \rightarrow 0$, we have

$$\liminf_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, n_k)] \geq - \sup_{x \in \mathbb{S}(\Gamma_k)} I_k(x). \quad (31)$$

Since $\Phi_m^*(z, x)$ is continuous in x , $\inf_{z \in A_k(d)} \Phi_m^*(z, x)$ is upper semi-continuous in x , and $I_k(x)$ is upper semi-continuous in x . In addition, $\mathbb{S}(\Gamma_k)$ is compact, therefore the supremum over x on the right hand side of (31) is a maximization. Consider the policy $x_k^* = \arg \max_{x \in \mathbb{S}(\Gamma_k)} I_k(x)$. For each $n_k \geq 1$, let agent k use the policy x_{k,n_k} where $x_{k,n_k}[\gamma] = \lfloor x_k^*[\gamma] n_k \rfloor / n_k$ for all $\gamma \in \Gamma_k$, and if $x_{k,n_k}[\gamma]$ do not sum to 1 over $\gamma \in \Gamma_k$, we simply choose the remaining private observations from an arbitrary distribution, and ignore them when making the decision for agent k . We have $x_{k,n_k} \rightarrow x_k^*$ as $n_k \rightarrow \infty$. We first show a simple lemma.

Lemma B.1. *Suppose that Assumption 2 holds, agent k adopts the policy $x_{k,n_k} = (x_{k,n_k}[\gamma])_{\gamma \in \Gamma_k}$ when it has access to n_k*

⁹Note that Assumption 3 is required in (28), without which we need to replace $\mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d, n_k) < 0)$ with $\max\{\mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d, n_k) < 0), \mathbb{P}_m(g_k(\bar{Z}_{n_k}^0(x_{k,n_k}), d, n_k) = 0)\}$ if expert k uses randomization to produce its final decision. The second term in the maximization unfortunately cannot be characterized using our existing approach.

private observations, and $x_{k,n_k} \rightarrow x_k$ as $n_k \rightarrow \infty$. Then, the sequence of random variables $(\bar{Z}_{n_k}(x_{k,n_k}))_{n_k \geq 1}$ defined in (10) satisfies a LDP under hypothesis $H = m$, for every $m \in [0, M-1]$, with good rate function $\Phi_m^*(\cdot, x_k)$.

Proof: We apply the Gärtner-Ellis Theorem [35] to prove the lemma. Let $Z_i = (\log \ell_{m0}(Y_k[i]))_{m=1}^{M-1}$. We have $n_k \bar{Z}_{n_k}(x_{k,n_k}) = \sum_{i=1}^{n_k} Z_i$ since $Z_i, i = 1 \dots, n_k$ are independent. For every $t \in \mathbb{R}^{M-1}$, we obtain

$$\begin{aligned} & \frac{1}{n_k} \log \mathbb{E}_m [\exp(\langle n_k t, \bar{Z}_{n_k}(x_{k,n_k}) \rangle)] \\ &= \frac{1}{n_k} \log \mathbb{E}_m \left[\exp \left(\left\langle t, \sum_{i=1}^{n_k} Z_i \right\rangle \right) \right] \\ &= \frac{1}{n_k} \sum_{i=1}^{n_k} \log \mathbb{E}_m [\exp(\langle t, Z_i \rangle)] \\ &= \sum_{\gamma \in \Gamma_k} x_{k,n_k}[\gamma] \log \mathbb{E}_m [\exp(\langle t, (\log \ell_{m0}^\gamma)_{m=1}^{M-1} \rangle)] \\ &\rightarrow \varphi_m(t, x_k), \end{aligned}$$

as $n_k \rightarrow \infty$. From Assumption 2 and Lemma 2.3.9 of [35], we have $\Phi_m^*(z, x_k)$ is a good rate function, and the lemma follows from the Gärtner-Ellis Theorem. ■

From Lemma 1, we have for each $d \in [0, d_k - 1]$,

$$\limsup_{n_k \rightarrow \infty} \sup_z |g_k(z, d, n_k) - f_k(z, d)| = 0,$$

and applying Theorem 4.2.23 of [35], we obtain from (28) and Lemma B.1 that

$$\limsup_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \leq - \inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*). \quad (32)$$

We then have

$$\begin{aligned} & \limsup_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, m_l)] \\ &\leq \limsup_{n_k \rightarrow \infty} \max_{\substack{0 \leq m \leq M-1 \\ 0 \leq d \leq d_k-1}} \left\{ \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) - c_k(m, d) \right\} \\ &\leq -I_k(x_k^*), \end{aligned} \quad (33)$$

where the last inequality follows from (32). Finally, (31) together with (33) gives us claim (i).

To show claim (ii), fix any $x_k^* \in \arg \max_{x \in \mathbb{S}(\Gamma_k)} I_k(x)$. We note that if there exists a subsequence of policies $(x_{k,m_l})_{l \geq 1}$ with $\lim_{l \rightarrow \infty} x_{k,m_l} = x_k$ and $I_k(x_k) < I_k(x_k^*)$, then using the same arguments that lead to (31), we have

$$\liminf_{n_k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}[C_k(H, D_k, n_k)] \geq -I_k(x_k) > -I_k(x_k^*),$$

a contradiction to (33). Therefore, each policy subsequence converges to some x_k with $I_k(x_k) = I_k(x_k^*)$, and there is no loss in optimality if we restrict the sequence of policies to converge to x_k^* .

Finally, to show claim (iii), we have from (30) that

$$\liminf_{l \rightarrow \infty} \frac{1}{n_k} \log \mathbb{P}_m(D_k = d) \geq - \inf_{z \in A_k(d)} \Phi_m^*(z, x_k^*), \quad (34)$$

since $\Phi_m^*(z, x)$ is continuous in x . Together with (32), the claim now follows, and the theorem is proved.

C. Proof of Lemma 3

The non-negativity and convexity of $\Phi_m(z, x)$ follows from Lemma 2.2.31 of [35]. From Jensen's inequality, for any $t \in \mathbb{R}^{M-1}$, we have

$$\varphi_m(t, x) \geq \sum_{\gamma} x[\gamma] \mathbb{E}_m [\langle t, Z^\gamma \rangle] = \langle t, \tilde{z}_m(x) \rangle,$$

which implies that $\Phi_m(\tilde{z}, x) = 0$, and the lemma is proved.

D. Proof of Theorem 4

We first present a generalization of Theorem 5 of [41] (see also [23] for a slightly more updated version). The proof steps are similar to that in [41] and [23], and are provided below for completeness.

Proposition B.1. *Suppose that Assumptions 1 and 2 hold, and agent 0 adopts the opinion of agent $k \geq 1$ and policy x_0 . Let*

$$s_{ij}^* = \operatorname{argmax}_{s \in [0,1]} \left\{ \frac{s}{n_0} \log \frac{C_0(i, j, n_0) \mathbb{P}_i(D_k = d)}{C_0(j, i, n_0) \mathbb{P}_j(D_k = d)} - \Lambda_{ij}(s, x_0) \right\}$$

where $\Lambda_{ij}(\cdot, \cdot)$ is as defined in (6). Then, for any $\epsilon > 0$, and any $d \in [0, d_k - 1]$, there exists n such that for all $n_0 \geq n$, we have for all $i \neq j$,

$$\begin{aligned} & \frac{1}{n_0} \log \left\{ \frac{\min_{j': j' \neq i} C_0(i, j', n_0) \mathbb{P}_i(D_0(k) \neq i, D_k = d)}{+ C_0(j, i, n_0) \mathbb{P}_j(D_0(k) = i, D_k = d)} \right\} \\ & \geq (1 - s_{ij}^*) \left(\frac{1}{n_0} \log \mathbb{P}_i(D_k = d) - c_0(i) \right) \\ & \quad + s_{ij}^* \left(\frac{1}{n_0} \log \mathbb{P}_j(D_k = d) - c_0(j) \right) + \Lambda_{ij}(s_{ij}^*, x_0) - \epsilon. \end{aligned} \quad (35)$$

Proof: From Theorem 5 of [41] (or Proposition A.2 of [23]), we have for $i, j \in [0, M-1]$, with $j \neq i$, and every $s \in [0, 1]$, either

$$\begin{aligned} & \mathbb{P}_i(D_0(k) \neq i \mid D_k = d) \\ & \geq \frac{1}{4} \exp \left(n_0 \Lambda_{ij}(s, x_0) - s n_0 \frac{d}{ds} \Lambda_{ij}(s, x_0) \right. \\ & \quad \left. - s \sqrt{2 n_0 \frac{d^2}{ds^2} \Lambda_{ij}(s, x_0)} \right), \end{aligned} \quad (36)$$

or

$$\begin{aligned} & \mathbb{P}_j(D_0(k) = i \mid D_k = d) \\ & \geq \frac{1}{4} \exp \left(n_0 \Lambda_{ij}(s, x_0) + (1 - s) n_0 \frac{d}{ds} \Lambda_{ij}(s, x_0) \right. \\ & \quad \left. - (1 - s) \sqrt{2 n_0 \frac{d^2}{ds^2} \Lambda_{ij}(s, x_0)} \right). \end{aligned} \quad (37)$$

If $s_{ij}^* \in (0, 1)$, we have

$$\frac{d}{ds} \Lambda_{ij}(s_{ij}^*, x_0) = \frac{1}{n_0} \log \frac{C_0(i, j, n_0) \mathbb{P}_i(D_k = d)}{C_0(j, i, n_0) \mathbb{P}_j(D_k = d)},$$

and using Lemma A.2, (36) and (37), we obtain

$$\begin{aligned} & \frac{1}{n_0} \log \left\{ \frac{\min_{j': j' \neq i} C_0(i, j', n_0) \mathbb{P}_i(D_0(k) \neq i, D_k = d)}{+ C_0(j, i, n_0) \mathbb{P}_j(D_0(k) = i, D_k = d)} \right\} \\ & \geq \frac{1 - s_{ij}^*}{n_0} \log \left(\min_{j': j' \neq i} C_0(i, j', n_0) \mathbb{P}_i(D_k = d) \right) \\ & \quad + \frac{s_{ij}^*}{n_0} \log (C_0(j, i, n_0) \mathbb{P}_j(D_k = d)) \\ & \quad + \Lambda_{ij}(s_{ij}^*, x_0) - \frac{1}{n_0} \log 2 - \sqrt{\frac{2G(2)}{n_0}} \\ & \geq (1 - s_{ij}^*) \left(\frac{1}{n_0} \log \mathbb{P}_i(D_k = d) - c_0(i) \right) \\ & \quad + s_{ij}^* \left(\frac{1}{n_0} \log \mathbb{P}_j(D_k = d) - c_0(j) \right) + \Lambda_{ij}(s_{ij}^*, x_0) - \epsilon, \end{aligned}$$

where the last inequality follows from Assumption 1 for n_0 sufficiently large. On the other hand, if $s_{ij}^* = 0$, we have

$$\frac{d}{ds} \Lambda_{ij}(0, x_0) \geq \frac{1}{n_0} \log \frac{C_0(i, j, n_0) \mathbb{P}_i(D_k = d)}{C_0(j, i, n_0) \mathbb{P}_j(D_k = d)},$$

since $\Lambda_{ij}(s, x_0)$ is convex s . The inequality (35) then holds trivially. A similar argument holds for $s_{ij}^* = 1$, and the proposition is proved. ■

We next proceed to prove Theorem 4. Suppose that agent 0 adopts the policy x_0 . For any $\epsilon > 0$, and for n_0 sufficiently large, we have

$$\begin{aligned} & \frac{1}{n_0} \log \mathbb{E} [C_0(H, D_0(k), n_0)] \\ & = \frac{1}{n_0} \log \sum_{d=0}^{d_k-1} \sum_{i \neq j} \pi_i C_0(i, j, n_0) \mathbb{P}_i(D_0(k) = j, D_k = d) \\ & \geq \max_{\substack{i \neq j \\ 0 \leq d \leq d_k-1}} \frac{1}{n_0} \log \left\{ \frac{\min_{j': j' \neq i} C_0(i, j', n_0) \mathbb{P}_i(D_0(k) \neq i, D_k = d)}{+ C_0(j, i, n_0) \mathbb{P}_j(D_0(k) = i, D_k = d)} \right\} \\ & \quad - \frac{1}{n_0} \log 2 + \min_m \frac{1}{n_0} \log \pi_m \\ & \geq \max_{\substack{i \neq j \\ 0 \leq d \leq d_k-1}} \min_{s \in [0,1]} \left\{ (1 - s) \left(\frac{1}{n_0} \log \mathbb{P}_i(D_k = d) - c_0(i) \right) \right. \\ & \quad \left. + s \left(\frac{1}{n_0} \log \mathbb{P}_j(D_k = d) - c_0(j) \right) + \Lambda_{ij}(s, x_0) \right\} - \epsilon, \end{aligned} \quad (38)$$

where the last inequality follows from Proposition B.1. By letting $n_0 \rightarrow \infty$ and $\epsilon \rightarrow 0$ in (38), we obtain

$$\begin{aligned} & \limsup_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{E} [C_0(H, D_0(k), n_0)] \\ & \geq \max_{\substack{i \neq j \\ 0 \leq d \leq d_k-1}} \min_{s \in [0,1]} \left\{ (1 - s) \left(\lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{P}_i(D_k = d) - c_0(i) \right) \right. \\ & \quad \left. + s \left(\lim_{n_0 \rightarrow \infty} \frac{1}{n_0} \log \mathbb{P}_j(D_k = d) - c_0(j) \right) + \Lambda_{ij}(s, x_0) \right\}, \end{aligned} \quad (39)$$

and the lower bound (16) follows from Theorem 2(iii).

We next show that there exists a decision rule for agent 0 that achieves $\mathcal{E}_0(x_0)$ in (16). Given $D_k = d$, consider the following rule to differentiate between hypotheses $H = i$

and $H = j$ for $i \neq j$: declare $H = i$ iff $\frac{1}{n_0} \log \ell_{ji}(Y_0[1 : n_0]) \leq h_{ji}$, where $h_{ji} = -q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) + q_k \inf_{z \in A_k(d)} \Phi_j^*(z, x_k^*) - c_0(i) + c_0(j)$.

By a simple generalization of Cramér's Theorem¹⁰ [35], and Theorem 2(iii), we have for every $\epsilon > 0$ and all n_0 sufficiently large,

$$\begin{aligned} & \frac{1}{n_0} \log(C_0(i, j, n_0) \mathbb{P}_i(D_0(k) = j, D_k = d)) \\ &= \frac{1}{n_0} \log C_0(i, j, n_0) + \frac{1}{n_0} \log \mathbb{P}_i(D_k = d) \\ &+ \frac{1}{n_0} \log \mathbb{P}_i(D_0(k) = j \mid D_k = d) \\ &\leq -c_0(i) - q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) \\ &- \max_{s \in [0, 1]} \{s h_{ji} - \Lambda_{ij}(s, x_0)\} + \epsilon \\ &\leq - \max_{s \in [0, 1]} \left\{ (1-s) \left(q_k \inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) + c_0(i) \right) \right. \\ &\quad \left. + s \left(q_k \inf_{z \in A_k(d)} \Phi_j^*(z, x_k^*) + c_0(j) \right) - \Lambda_{ij}(s, x_0) \right\} + \epsilon, \end{aligned}$$

from which we obtain

$$\begin{aligned} & \frac{1}{n_0} \log \mathbb{E}[C_0(H, D_0(k), n_0)] \\ &\leq \max_{\substack{i \neq j \\ 0 \leq d \leq d_k - 1}} \frac{1}{n_0} \log(C_0(i, j, n_0) \mathbb{P}_i(D_0(k) = j, D_k = d)) \\ &+ \frac{1}{n_0} \log M \\ &\leq -\mathcal{E}_0(k, x_0) + \epsilon. \end{aligned}$$

By taking $n_0 \rightarrow \infty$ and $\epsilon \rightarrow 0$, we obtain the theorem by maximizing $\mathcal{E}_0(x_0)$ over all policies x_0 . The proof is now complete.

E. Proof of Theorem 5

The first part of the theorem is a direct consequence of Theorem 4. From Lemma 3, we have $\inf_{z \in A_k(d)} \Phi_i^*(z, x_k^*) \geq 0$ for any $i \in [0, M-1]$, $k \geq 1$, $d \in [0, d_k-1]$, and policy x_k^* . Furthermore, Assumption 1 implies that $c_0(i) \geq 0$ for all $i \in [0, M-1]$. Therefore, the inequalities (20) and (21) follow from (16), (17), (18), and (19), and the proof is complete.

F. Proof of Proposition 6

Suppose that agent k adopts the policy x_k . From the pigeonhole principle, if $d_k < M$, there exists a region $A_k(d)$ in which both $\Phi_i(\cdot, x_k)$ and $\Phi_j(\cdot, x_k)$ achieve their minimum value of 0, for some $i \neq j$. Since for any (i', j') , we have

$$\begin{aligned} & \max_{s \in [0, 1]} \left\{ (1-s) \left(q_k \inf_{z \in A_k(d)} \Phi_{i'}^*(z, x_k^*) + c_0(i') \right) \right. \\ &\quad \left. + s \left(q_k \inf_{z \in A_k(d)} \Phi_{j'}^*(z, x_k^*) + c_0(j') \right) - \Lambda_{i'j'}(s, x_0) \right\} \\ &\geq \max_{s \in [0, 1]} (-\bar{\Lambda}_{i'j'}(s, x_0, c_0)) \\ &= - \min_{s \in [0, 1]} \bar{\Lambda}_{ij}(s, x_0, c_0), \end{aligned}$$

¹⁰Cramér's Theorem applies to independent and identically distributed (i.i.d.) observations. The private observations $Y_0[1 : n_0]$ are not i.i.d., but are independent and can be divided into groups of i.i.d. observations.

we obtain from (16) and (17),

$$\mathcal{E}_0(k, x_0) = - \min_{s \in [0, 1]} \bar{\Lambda}_{ij}(s, x_0, c_0) = \mathcal{E}_{0,B}(k, x_0),$$

and the proposition is proved.

G. Proof of Proposition 7

Since the proof is similar to that of Theorem 4, we provide only an outline here. From the proposition assumptions, we have for every $j \in [0, M-1]$, $\mathbb{P}_j(D_k = j)$ is bounded away from zero, i.e., $\lim_{n_k \rightarrow \infty} (1/n_k) \log \mathbb{P}_j(D_k = j) = 0$ because otherwise the expected loss of agent k can be decreased. Therefore, we have $\inf_{z \in A_k(j)} \Phi_j^*(z, x_k) = 0$. By comparing (17) (with x_0 replaced by x_k) and (14), and using an inductive argument, we have $\inf_{z \in A_k(i)} \Phi_j^*(z, x_k) = \Lambda_{ji}^*(c_k(i) - c_k(j), x_k)$, and (23) follows.

We next show the second part of the proposition. Suppose that for some $i \neq j$, we have $c_0(i) = c_0(j)$ and (24) holds. Then from (22), we have $\bar{\Lambda}_{ij}(s, x_0, c_0) = \bar{\Lambda}_{ji}(s, x_0, c_0)$ for all $s \in [0, 1]$. Let $g_{ij}(\Delta) = \max_{s \in [0, 1]} \{s q_k \Lambda_{ji}^*(\Delta, x_k) - \bar{\Lambda}_{ij}(s, x_0, c_0)\}$. Note that $g_{ij}(\Delta)$ is non-decreasing in Δ . Therefore, since $\bar{\Lambda}_{ij}(s, x_0, c_0)$ is symmetrical on $s \in [0, 1]$, we have $\min(g_{ij}(\Delta), g_{ji}(-\Delta))$ is maximized if $\Lambda_{ji}^*(\Delta, x_k) = \Lambda_{ij}^*(-\Delta, x_k)$, which holds if $\Delta = 0$. Take $\Delta = c_k(i) - c_k(j)$, and the proposition follows.

APPENDIX C

CHARACTERIZATION OF ASYMPTOTIC DECISION REGIONS

In this appendix, we give a characterization for the asymptotic decision region $A_k(d)$ for an agent k , and $d \in [0, d_k-1]$. For $i, j \in [0, M-1]$ and $p, q \in [0, d_k]$, define the halfspace

$$\begin{aligned} B_k(i, p, j, q) &= \{z = (z[m])_{1 \leq m \leq M-1} \in \mathbb{R}^{M-1} : \\ &\quad z[i] - z[j] \geq c_k(i, p) - c_k(j, q)\}, \end{aligned} \quad (40)$$

where $z[0] = 0$. For each $p \in [0, d_k-1]$, let $m_p \in [0, M-1]$ be a chosen corresponding index. From (12), we have $z \in \cap_{i \neq m} B_k(m_p, p, i, p)$ iff $\tilde{f}_k(z, p) = z[m_p] - c_k(m_p, p)$.

Consider a $z \in \mathbb{R}^{M-1}$ such that $f(z, d) < 0$. Then, there exists a sequence $(m_p)_{0 \leq p \leq d_k-1} \in [0, M-1]^{d_k}$ such that $\tilde{f}_k(z, p) = z[m_p] - c_k(m_p, p)$ for all $p \in [0, d_k-1]$ and $\tilde{f}_k(z, p) - \tilde{f}_k(z, d) = z[m_p] - z[m_d] - c_k(m_p, p) + c_k(m_d, d) > 0$ for all $p \neq d$, i.e.,

$$\begin{aligned} z &\in H_d((m_p)_{p=0}^{d_k-1}) \\ &\triangleq \bigcap_{p=0}^{d_k-1} \bigcap_{i \neq m_p} B_k(m_p, p, i, p) \bigcap_{p \neq d} B_k(m_p, p, m_d, d), \end{aligned}$$

where $H_d((m_p)_{p=0}^{d_k-1})$ is a polyhedron since it consists of intersections of halfspaces. On the other hand, if such a sequence $(m_p)_{p=0}^{d_k-1}$ exists, then $z \in A_k(d)$. Therefore, the set $A_k(d)$ is the union over all sequences $(m_p)_{p=0}^{d_k-1} \in [0, M-1]^{d_k}$ of the polyhedra $H_d((m_p)_{p=0}^{d_k-1})$.

REFERENCES

- [1] S. Aral and D. Walker, "Identifying influential and susceptible members of social networks," *Science*, vol. 337, no. 6092, pp. 337–341, 2012.
- [2] A. Mitchell, J. Kiley, J. Gottfried, and E. Guskin. (2013, Oct.) The role of news on Facebook. Pew Research Center. [Online]. Available: http://www.journalism.org/files/2013/10/facebook_news_10-24-2013.pdf
- [3] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer: quantifying influence on Twitter," in *Proc. 4th ACM International Conference on Web Search and Data Mining*, 2011.
- [4] W. Luo, W. P. Tay, and M. Leng, "Identifying infection sources and regions in large networks," *IEEE Trans. Signal Process.*, vol. 61, no. 11, pp. 2850–2865, Jun. 2013.
- [5] J. Leskovec, L. A. Adamic, and B. A. Huberman, "The dynamics of viral marketing," *ACM Trans. Web*, vol. 1, May 2007.
- [6] Y. Zhang, Z. Wang, and C. Xia, "Identifying key users for targeted marketing by mining online social network," in *Proc. IEEE International Conference on Advanced Information Networking and Applications Workshops*, 2010.
- [7] Zero Hedge. [Online]. Available: <http://www.zerohedge.com/>
- [8] DealBook. [Online]. Available: <http://dealbook.nytimes.com/>
- [9] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, Mar. 2011.
- [10] Y. Kanoria and O. Tamuz, "Tractable Bayesian social learning on trees," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 4, pp. 756–765, April 2013.
- [11] G. Kortuem, F. Kawsar, D. Fitton, and V. Sundramoorthy, "Smart objects as building blocks for the Internet of Things," *IEEE Internet Computing*, vol. 14, no. 1, pp. 44–51, 2010.
- [12] Y. Ding, Y. Jin, L. Ren, and K. Hao, "An intelligent self-organization scheme for the Internet of Things," *IEEE Computational Intelligence Magazine*, vol. 8, no. 3, pp. 41–53, 2013.
- [13] J. N. Tsitsiklis, "Extremal properties of likelihood-ratio quantizers," *IEEE Trans. Commun.*, vol. 41, no. 4, pp. 550–558, 1993.
- [14] B. Chen and P. K. Varshney, "A Bayesian sampling approach to decision fusion using hierarchical models," *IEEE Trans. Signal Process.*, vol. 50, no. 8, pp. 1809–1818, Aug. 2002.
- [15] J.-F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 407–416, 2003.
- [16] Y. Lin, B. Chen, and P. K. Varshney, "Decision fusion rules in multi-hop wireless sensor networks," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 2, pp. 475–488, Apr. 2005.
- [17] P. Willett, P. Swaszek, and R. Blum, "The good, bad and ugly: distributed detection of a known signal in dependent Gaussian noise," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3266 – 3279, Dec. 2000.
- [18] H. Chen, B. Chen, and P. Varshney, "A new framework for distributed detection with conditionally dependent observations," *IEEE Trans. Signal Process.*, vol. 60, no. 3, pp. 1409 –1419, Mar. 2012.
- [19] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "Data fusion trees for detection: Does architecture matter?" *IEEE Trans. Inf. Theory*, vol. 54, no. 9, pp. 4155–4168, Sep. 2008.
- [20] —, "On the impact of node failures and unreliable communications in dense sensor networks," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2535–2546, Jun. 2008.
- [21] —, "Bayesian detection in bounded height tree networks," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 4042–4051, Oct. 2009.
- [22] Z. Zhang, A. Pezeshki, W. Moran, S. Howard, and E. K. P. Chong, "Error probability bounds for balanced binary relay trees," *IEEE Trans. Inf. Theory*, vol. 58, no. 6, pp. 3548–3563, 2012.
- [23] W. P. Tay, "The value of feedback in decentralized detection," *IEEE Trans. Inf. Theory*, vol. 58, no. 12, pp. 7226–7239, Dec. 2012.
- [24] J. N. Tsitsiklis, "Decentralized detection," *Advances in Statistical Signal Processing*, vol. 2, pp. 297–344, 1993.
- [25] J. D. Papastavrou and M. Athans, "Distributed detection by a large team of sensors in tandem," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 639–653, 1992.
- [26] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "On the sub-exponential decay of detection error probabilities in long tandems," *IEEE Trans. Inf. Theory*, vol. 54, no. 10, pp. 4767–4771, Oct. 2008.
- [27] K. Drakopoulos, A. Ozdaglar, and J. Tsitsiklis, "On learning with finite memory," *IEEE Trans. Inf. Theory*, vol. 59, no. 10, pp. 6859–6872, Oct. 2013.
- [28] J. N. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Trans. Autom. Control*, vol. 30, pp. 440–446, 1985.
- [29] S. Nitinawarat, G. Atia, and V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.
- [30] E. Tuncel, "On error exponents in hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 51, no. 8, pp. 2945–2950, Aug. 2005.
- [31] N. Grigoryan and A. Harutyunyan, "Error exponents in multiple hypothesis testing for arbitrarily varying sources," in *IEEE Information Theory Workshop*, Aug. 2010, pp. 1–5.
- [32] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals, Syst.*, vol. 1, pp. 167–182, 1988.
- [33] T. A. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York, NY: John Wiley & Sons, Inc., 2005.
- [34] J. Ho, W. P. Tay, and T. Q. S. Quek, "Robust detection and social learning in tandem networks," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Florence, Italy, May 2014.
- [35] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. New York, NY: Springer-Verlag, 1998.
- [36] H. P. Young, "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning," *American Economic Review*, vol. 99, no. 5, pp. 1899–1924, 2009. [Online]. Available: <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.1899>
- [37] X. Vives, "Social learning and rational expectations," *European Economic Review*, vol. 40, no. 3–5, pp. 589 – 601, 1996, papers and Proceedings of the Tenth Annual Congress of the European Economic Association. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0014292195000720>
- [38] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [39] V. Veeravalli, T. Basar, and H. Poor, "Minimax robust decentralized detection," in *IEEE Trans. on Inf. Theory*, vol. 40, 1994, pp. 35–40.
- [40] D. Stroock, *Probability Theory: An Analytic View*. Cambridge, UK: Cambridge University Press, 1993.
- [41] C. E. Shannon, R. G. Gallager, and E. R. Berlekamp, "Lower bounds to error probability for coding on discrete memoryless channels, I," *Information and Control*, vol. 10, pp. 65–103, 1967.



Wee Peng Tay (S'06 M'08) received the B.S. degree in Electrical Engineering and Mathematics, and the M.S. degree in Electrical Engineering from Stanford University, Stanford, CA, USA, in 2002. He received the Ph.D. degree in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2008. He is currently an Assistant Professor in the School of Electrical and Electronic Engineering at Nanyang Technological University, Singapore. His research interests include distributed detection and

estimation, distributed signal processing, sensor networks, social networks, information theory, and applied probability.

Dr. Tay received the Singapore Technologies Scholarship in 1998, the Stanford University President's Award in 1999, and the Frederick Emmons Terman Engineering Scholastic Award in 2002. He is the coauthor of the best student paper award at the 46th Asilomar conference on Signals, Systems, and Computers. He is currently serving as the chair of DSNIG in IEEE MMTC, and has served as a technical program committee member for various international conferences.